DOCUMENT RESUME

ED 153 722 PS 009 682

AUTHOR Moran, James D., III; McCullers, John C.
TITLE Reward and Number of Choices in Children's Probability Learning: An Attempt to Reconcile Conflicting Findings.
SPONS AGENCY National Inst. of Mental Health (DHEW), Rockville, Md.
PUB DATE Mar 78
GRANT NIMH-MH-22041; NIMH-MH-26359
NOTE 19p.; Paper presented at the Southwestern Society for Research in Human Development Conference (Dallas, Texas, March 17-18, 1978); Paper is based on master's thesis of first author (University of Oklahoma)

EDRS PRICE MF-$0.83 HC-$1.67 Plus Postage.
DESCRIPTORS Age Differences; *College Students; Conflict Resolution; *Educational Research; *Elementary School Students; Expectation; Learning Theories; Models; Multiple Choice Tests; *Probability Theory; Response Style (Tests); *Rewards; Sex Differences; *Task Performance; Visual Stimuli

ABSTRACT

This paper reports on a study conducted to clarify issues involving reward effects in children's probability learning; specifically, the detrimental-effects-of-reward theory, and the utility and expectancy theories. Five event schedules in which the probabilistic versus partial reinforcement dimension was separated from the number of choices in the task were examined in a size discrimination task at two levels of reward (pennies and plastic token markers). The subjects were 100 fourth graders and 100 college students of both sexes. Whether or not the event probabilities summed to unity proved to be a significantly more important determiner of terminal performance than the number of response alternatives in the task. Females maximized significantly more than males and college students significantly more than fourth graders. Significantly greater maximization was also exhibited by fourth graders who were in the high reward condition. These reward effects appeared to be more consistent with a detrimental-effects-of-reward interpretation than either utility or expectancy interpretations. (Author/CA)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

ED153722

Reward and Number of Choices in Children's Probability Learning: An Attempt to Reconcile Conflicting Findings

James D. Moran III and John C. McCullers

Oklahoma State University

Preprint: This paper is under consideration for publication. Any quotation of this material is not allowed without expressed written permission from the author.

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

James D. Moran III

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC) AND USERS OF THE ERIC SYSTEM."

This paper is based on the master's thesis of the first author conducted under the direction of the second author and submitted to the graduate college of the University of Oklahoma. This research was supported in part by NIMH grants MH 22041 and MH 26359 to the second author. The authors wish to thank Lester Reed, Superintendent of Norman Public Schools and the principals of Eisenhower and John F. Kennedy Schools for their cooperation and assistance in making children available for the study.

The current address of both authors is now Department of Family Relations and Child Development, Oklahoma State University, Stillwater, Oklahoma, 74074. Requests for reprints should be sent to the first author.

PS 009882

Running head: Reward and Number of Choices in Probability Learning

Abstract

Five event schedules (75:0, 75:25, 75:0:0, 75:25:0, 75:12.5:12.5) were examined in a size discrimination task at two levels of reward (pennies and plastic token markers). The subjects were 100 fourth graders and 100 college students of both sexes. Whether or not the event probabilities summed to unity proved to be a more important determiner of terminal performance than the number of response alternatives in the task ($p < .01$). Females maximized more than males ($p < .05$) and college students more than fourth graders ($p < .001$). Greater maximization was also exhibited by fourth graders who were in the high reward condition ($p < .01$). These reward effects appeared to be more consistent with a detrimental-effects-of-reward interpretation than either utility or expectancy interpretations.

Reward and Number of Choices in Children's Probability Learning: An Attempt to Reconcile Conflicting Findings

Investigations involving reward effects in children's probability learning have produced conflicting results. In some studies (Brackbill, Kappy, & Starr, 1962; Siegel & Andrews, 1962) greater maximization occurred with higher levels of reward. In others (e.g., Stevenson and Weir, 1959) lower reward levels led to more maximizing.

One explanation of these differences has been in terms of the number of choice alternatives in the task. In some studies (Weir & Gruen, 1965) terminal performance revealed greater maximization on a two-choice task than on a three-choice task. These findings, however, run counter to the results of Gardner (1957, 1958) who found that as the number of choices increased so did the number of responses to the most frequently reinforced stimulus. The two- vs. three-choice controversy led Stevenson (1970) to conclude that, "the results of these studies are tantalizing, for although the effects are replicable, the bases of the differences in results remain unknown" (p. 898).

Interestingly, two- and three-choice tasks have typically differed along another potentially important dimension: the sum of the probability event schedules. Although there have been numerous variations in probability tasks, the three-choice task has been typically used by Stevenson and his colleagues (Stevenson & Hoving, 1964; Stevenson & Weir, 1959) and can generally be classified as a discrimination learning task involving partial reinforcement, e.g., 66:0:0 event schedules. The two-choice tasks used by other researchers (Brackbill et al., 1962; Siegel & Andrews, 1962) typically represent true probabilistic tasks in which the event schedules sum to unity,

e.g., 75:25.

Discussions of the probability literature have generally ignored the probabilistic vs. partial reinforcement aspect of the task. Reinforcement of more than one stimulus changes the nature of the task and perhaps the type of strategy that would be most effective. With a probabilistic schedule, the subject will receive some reinforcement for selecting the less frequently rewarded stimulus. Thus, when the subject chooses to maximize, he or she must forego that reinforcement and settle for less than 100% reinforcement. The subject in the partial reinforcement procedure faces a lesser dilemma in opting for maximization in tasks in which only one stimulus event is ever rewarded.

A related issue of theoretical interest is whether the tendency of older children and adults to maximize is better explained in terms of utility, the detrimental effects of reward, or expectancy. From the utility perspective (Brackbill et al., 1962; Offenbach, 1964) maximizing represents the most efficient approach to the problem in that it yields the greatest amount of reinforcement. From the perspective of the detrimental effects of reward, reward may either distract the subject's attention (Spence & Dunton, 1967; Spence & Segner, 1967) or produce a mechanical and stereotyped way of responding (McGraw & McCullers, 1974). In either case, the detrimental effect of reward would interfere with the subject's search for a perfect solution.

Differences between these two theoretical alternatives may be discernible in the rate of responding to the high probability event. The detrimental-effects-of-reward model would predict performance differences due to reward from the outset, that is, a child who is distracted by reward should be

affected very early in training. The utility model, however, should predict that the child initially would search for a perfect solution and finally adopt the stategy of maximizing, later in training, as a workable but second-best solution.

Expectancy theory provides a third explanatory possibility. The higher the subject's initial expectancy for success, the less likely he or she would be to opt for a maximizing strategy that would preclude a perfect solution and 100% reinforcement. We have no basis for estimating a subject's initial expectancy for success, of course, but there is no reason to believe that randomly grouped rewarded and nonrewarded subjects should differ systematically in initial expectancies.

To help clarify these issues, this study separated the probabilistic vs. partial reinforcement dimension from the number of choices in the task. Five event schedules were used. It was expected that performance across event schedules would differ mainly in terms of the number of choices at the outset of training but that the probabilistic vs. partial reinforcement dimension would prove to be a critical determiner of the terminal level of response. This prediction follows from Weir's (Weir, 1972; Wittig & Weir, 1971) suggestion that the difference in the frequency of occurrence between the high probability event and the other alternatives will determine the extent of maximizing at the terminal level of response. In the partial reinforcement schedules this difference in percentage payoffs is maximized.

Method

Subjects and design. The subjects were 100 fourth-grade, public school children and 100 university undergraduates. The university students' participation satisfied a requirement for the basic psychology course. The children and adults were predominately white, middle-class individuals.

Ten subjects (5 males and 5 females) were randomly assigned to each experimental condition with the exception that the fourth graders who were to receive tokens came from one school and those who were to receive monetary rewards came from a second school. This was done to prevent the subjects from entering the experiment with misconceptions as to what they would receive for participation.

Each subject received one of five probability schedules with either two choices (75:0, 75:25) or three choices (75:0:0, 75:25:0, 75:12.5:12.5) and one of two levels of rewards. Age and sex were included as additional factors to yield a 5 X 2 X 2 X 2 factorial design replicated across 5 blocks of 16 trials.

Materials. The stimuli were three black squares (9 cm^2 , 16 cm^2 , 25 cm^2) mounted on a 110.25 cm^2 white background. These were presented via a modified Wisconsin General Test Apparatus with three reward wells cut into the sliding tray. The stimulus cards were placed over the wells and held in place with plastic clamps. In the conditions that used only two choices, the third well was covered with a plain piece of wood and only the 16 cm^2 and 25 cm^2 stimuli were used.

To make a "choice", the subject lifted the stimulus card to expose the reward well. Reinforcers consisted of either pennies or plastic tokens

(for which plastic poker chips were used). The subject stacked the pennies or tokens in a container placed next to him.

Procedure. The subjects participated individually. They were informed that this was a learning experiment and that "a penny (plastic token) would be placed under the correct card". Each subject was told to try to "collect as many pennies (plastic tokens) as possible". Instructions were left vague as to whether or not a perfect solution was possible. The subjects in the reward conditions were also told that they could keep all the pennies that they collected. The subjects in the nonreward groups were not allowed to keep the plastic tokens, which served only as markers to identify correct responses.

The stimuli were presented simultaneously through the sliding door. The subject lifted up the card and either collected the reward or token from the well or, in the event of nonreinforcement, simply replaced the card on the tray. The presence or absence of a penny or token served as the only feedback as to the correctness of the subject's response and the task was self-paced. Following each trial the tray was brought back behind the sliding door and the experimenter reloaded the wells and rearranged the stimulus cards according to a preset, trial-defined schedule for eighty trials.

Reinforcement was randomized over blocks of 16 trials within the restrictions imposed by the five schedules. The position of each stimulus item was varied randomly with the restriction that each position (left, middle, right) was reinforced equally often.

Results and Discussion

An analysis of variance that followed the basic design of the study was performed using as scores the number of responses to the most frequently reinforced stimulus in a block of 16 trials. An abbreviated summary of this analysis, showing only the significant effects, is presented in Table 1.

Insert Table 1 about here

Age and sex effects. Overall, college students exhibited a greater tendency to maximize than fourth graders. Pairwise comparisons[1] confirmed that this difference was present on the first block of trials, $t(160) = 5.28$, $p < .01$, and continued through the last block of trials, $t(160) = 10.76$, $p < .01$. For both ages, responses to the high probability event increased across trials, but at a faster rate in the college students.

These developmental differences are in one sense in accordance with the traditional finding that college students maximize more than fourth graders (Weir, 1964). This difference, however, existed from the outset of training which does not support that notion that the performance of adults is initially similar to that of school-age children before moving to maximization as the best strategy (Goodwin, reported in Goulet & Goodwin, 1970; Jones & Liverant, 1960).

Although sex differences have been notably absent from the probability learning literature, the present study revealed that females did have a higher level of responding than males to the most frequently reinforced stimulus. This difference was present at both age levels but reached significance via individual comparison analysis only in the case of the fourth

graders, and then only on the last block of trials, $t(160) = 3.83$, $p < .05$.

Task factors. As predicted, the probabilistic vs. partial reinforcement dimension proved to be more important than number of choices in distinguishing between the event schedules at the terminal level of performance (see Figure 1). Scheffe's test of individual comparisons showed that on the first block of trials the two- vs. three-choice difference was significant $F(4,160) = 30.52$, $p < .01$, with a higher level of response in the two-choice conditions; the probabilistic vs. partial reinforcement dimension, on the other hand, failed to adequately separate the event schedules. However, by the end of training the importance of these factors had reversed. The two-choice tasks were no longer significantly different from the three-choice tasks, $F < 1$, but performance on the partial reinforcement tasks was significantly higher than on the tasks with event schedules that summed to unity, $F(4,160) = 14.28$, $p < .01$.

Insert Figure 1 about here

Weir (1972) has suggested that the choice of the most frequently reinforced stimulus should increase as the contrast in percentage payoff with the other alternatives increases. According to this view, responses to the high probability event should be low at the outset of training but increase over trials as the subject accumulates information about the payoff percentages. At the terminal-performance level the percentage payoff differences should determine the order of the event schedules. Weir's interpretation would account for the present schedules X trials interaction and the increasing separation, with training, of the unity schedules from those involving partial

reinforcement. It is within the probabilistic vs. partial reinforcement dimension that the differences in percentage payoffs are maximized.

Reward effects. Consistent with the findings of other researchers, (e.g., Stevenson and Hoving, 1964), no differences in performance were found for college students as a function of reward. However, with the fourth graders high reward did lead to increased selection of the high probability event, $t(160) = 4.84$, $p < .01$ (see Figure 2). In addition more fourth graders that received reward ($n = 15$) could be characterized as maximizers than fourth graders that received token feedback ($n = 5$), $\chi^2(1) = 5.06$, $p < .05$. Maximization, in this case, was defined as responding to the most frequently reinforced stimulus on at least 14 of the last 16 trials, a 87.5% response rate.

Insert Figure 2 about here

Both the utility and detrimental effects-of-reward hypotheses would predict greater maximization in the rewarded group. An essential distinction between these two theoretical viewpoints, however, was expected to be found in the rate of maximization. From the perspective of the detrimental effects of reward, a perfect solution would be sacrificed not in the interest of an optimizing strategy but rather as a result of centering on the reward throughout the task. The utility theory, however, predicts that the child will opt for the maximal payoff solution and thus begin by searching for a perfect solution, eventually arriving at maximization as the optimum strategy. The detrimental effects hypothesis suggests that differences between the rewarded and nonrewarded subjects exist from the outset of training, whereas the utility notion suggests that these differences develop during training.

Expectancy theory also would predict no differences between the reward-nonreward groups early in training since no pre-experimental differences between the groups are assumed to exist.

On the first block of trials, a nonsignificant trend, $t(160) = 1.09$, was found for a greater number of responses to be made to the most frequently reinforced stimulus for fourth graders under reward conditions. This trend reached statistical significance by the second block of trials, $t(160) = 3.05$, $p < .05$, and remained significant throughout the remaining blocks of trials, e.g., $t(160) = 4.28$, $p < .01$, on the last block. This rapid rate of maximization appears more consistent with a detrimental-effects-of-reward view than either utility or expectancy explanations.

Conclusion. The hypothesis that the probabilistic vs. partial reinforcement distinction rather than the number of choices would be a better predictor of terminal performance levels was confirmed. These results suggest the need to separate the variable of probabilistic vs. partial reinforcement from the number of choice alternatives in order to better understand the determinants of the child's performance in probability learning situations.

References

Brackbill, Y., Kappy, M. S., & Starr, R. H. Magnitude of reward and and probability learning. Journal of Experimental Psychology, 1962, 63, 32-35.

Gardner, R. A. Probability learning with two and three choices. American Journal of Psychology, 1957, 70, 174-185.

Gardner, R. A. Multiple-choice decision-behavior. American Journal of Psychology, 1958, 71, 710-717.

Goulet, L. R., & Goodwin, K. S. Development and choice behavior in probabilistic and problem-solving tasks. In H. W. Reese & L. P. Lipsitt (Eds.), Advances in child development and behavior (Vol. 5). New York: Academic Press, 1970.

Jones, M. H., & Liverant, S. Effects of age differences on choice behavior. Child Development, 1960, 31, 673-680.

McGraw, K. O., & McCullers, J. C. The distracting effect of material reward: An alternative explanation for the superior performance of reward groups in probability learning. Journal of Experimental Child Psychology, 1974, 18, 149-158.

Offenbach, S. I. Studies of children's probability learning behavior: I. Effect of reward and punishment at two age levels. Child Development, 1964, 35, 709-715.

Siegel, S., & Andrews, J. M. Magnitude of reinforcement and choice behavior in children. Journal of Experimental Psychology, 1962, 63, 337-341.

Spence, J. T., & Dunton, M. C. The influence of verbal and nonverbal reinforcement combinations in the discrimination learning of middle- and lower-class preschool children. Child Development, 1967, 38, 1177-1186.

Spence, J. T., & Segner, L. L. Verbal versus nonverbal reinforcement combinations in the discrimination learning of middle- and lower-class children. Child Development, 1967, 38, 29-38.

Stevenson, H. W. Learning in children. In P. H. Mussen (Ed.), Carmichael's manual of child psychology (Vol. 1). New York: John Wiley & Sons, 1970.

Stevenson, H. W., & Hoving, K. L. Probability learning as a function of age and incentive. Journal of Experimental Child Psychology, 1964, 1, 64-70.

Stevenson, H. W., & Weir, M. W. Variables affecting children's performance in a probability learning task. Journal of Experimental Psychology, 1959, 57, 403-412.

Weir, M. W. Developmental changes in problem-solving strategies. Psychological Review, 1964, 71, 473-490.

Weir, M. W. Probability performance: Reinforcement procedure and number of alternatives. American Journal of Psychology, 1972, 85, 261-270.

Weir, M. W., & Gruen, G. E. Role of incentive level, number of choice and type of task in children's probability learning, Journal of Experimental Child Psychology, 1965, 2, 121-134.

Wittig, M. A., & Weir, M. W. The role of reinforcement procedure in children's probability learning as a function of age and number of response alternatives. Journal of Experimental Child Psychology, 1971, 12, 228-239.

Footnote

[1]All individual comparisons and reported *t* tests refer to Tukey tests using pooled error terms.

TABLE 1

Analysis of Variance of Number of Responses to Most Frequently Reinforced Stimulus

Summary of Significant Effects

Source	df	MS	F-Ratio
Age	1	1452.22	45.74***
Sex	1	139.40	4.39*
Reward	1	124.86	2.93*
Schedule	4	131.37	4.14**
Age X Reward	1	167.44	5.27*
Error (Between)	160	31.75	
Trials[a]	4	265.48	60.29***
Age X Trials	4	22.35	5.08***
Reward X Trials	4	12.33	2.80*
Schedule X Trials	16	15.03	3.41***
Error (Within)	160	4.40	

[a]The Giesser-Greenhouse correction for violation of assumptions was also used for the within factor analysis. This conservative test yielded the following p- values: trials, $p < .01$; age X trials, $p < .05$; reward X trials, $p < .10$; schedule X trials, $p < .05$. No other values reported above were affected.

*$p < .05$

**$p < .01$

***$p < .001$

Figure Captions

Figure 1. Mean numbers of responses to the most frequently reinforced stimulus in each of the five event schedules. Filled circles denote partial reinforcement schedules and open circles and triangles denote probabilistic schedules; solid lines represent two-choice and broken lines represent three-choice tasks.

Figure 2. Mean number of responses to the most frequently reinforced stimulus as a function of age and level of reward. High reward-subjects received a penny for each correct response, and kept all money earned in the task. The low-reward subjects received a plastic token for each correct response, and the tokens served only as a source of feedback.

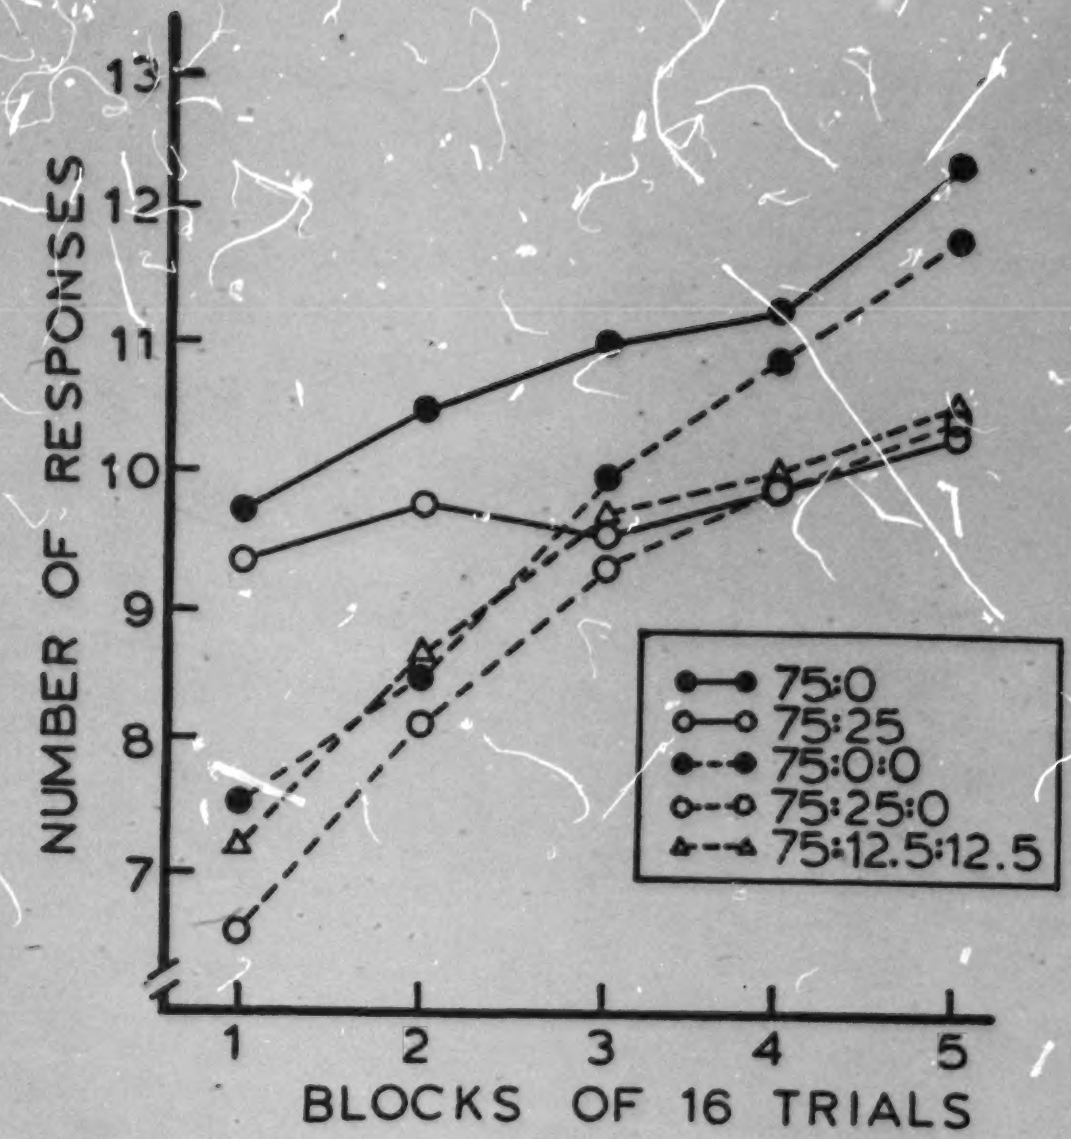
NUMBER OF RESPONSES
13
12
11
10
9
8
7
1
2
3
4
5
BLOCKS OF 16 TRIALS
75:0
75:25
75:0:0
75:25:0
75:12.5:12.5

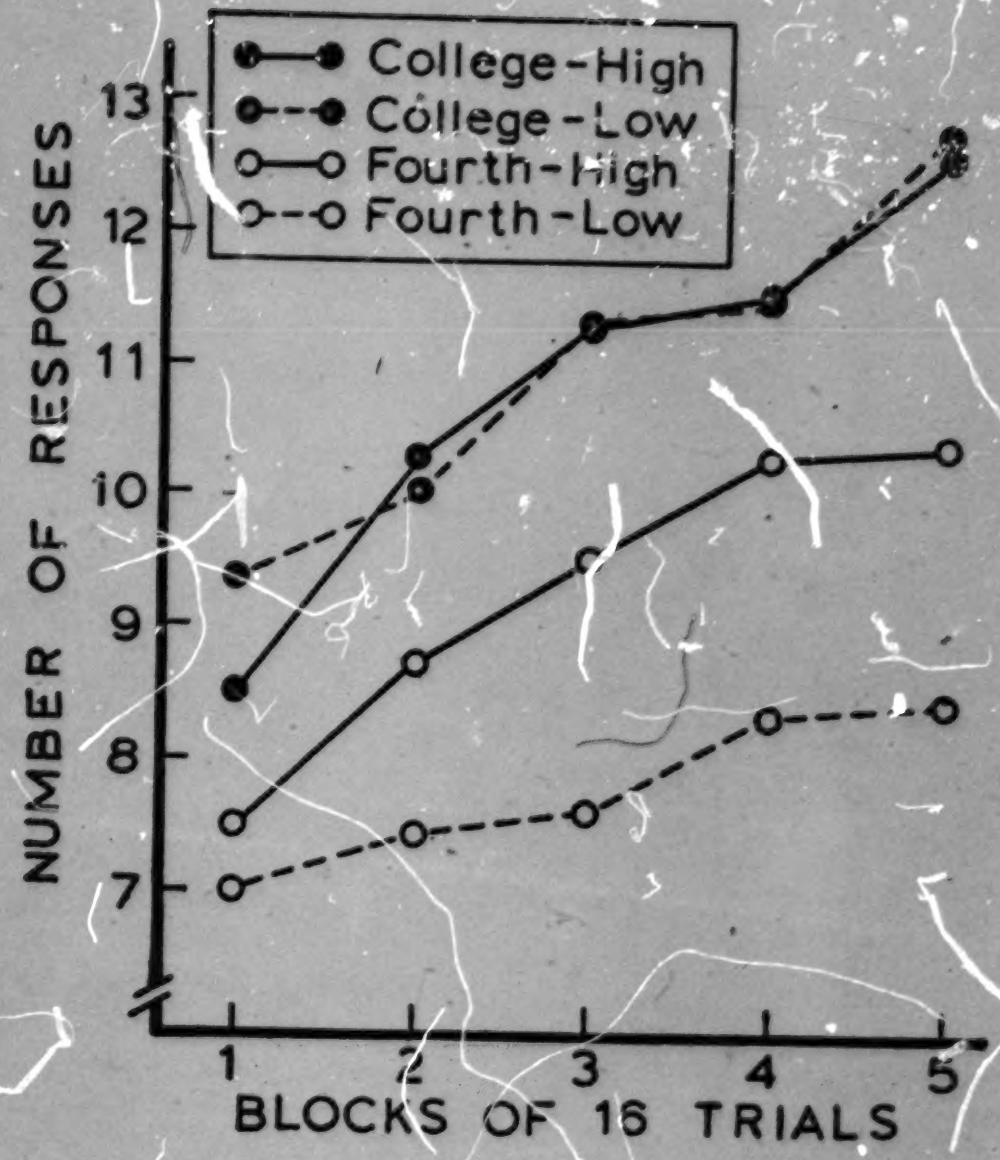
College-High
College-Low
Fourth-High
Fourth-Low
NUMBER OF RESPONSES
13
12
11
10
9
8
7
1
2
3
4
5
BLOCKS OF 16 TRIALS